The KANT Perspective: A Critique of Pure Transfer (and Pure Interlingua, Pure Statistics, ...)

Jaime G. Carbonell
Teruko Mitamura
Eric H. Nyberg, 3rd
Center for Machine Translation
Carnegie Mellon University
Pittsburgh, PA 15213

Abstract

There is a strong tendency among MT researchers to emphasize paradigmatic differences in MT approaches. This often leads to energetic criticism of competing systems, but obscures the fact that many techniques developed within a particular paradigm can improve the quality of MT systems in general. In this paper we show how practical MT development must move beyond dogmatic dismissal of differing approaches to an integrated, rational approach to MT which combines the best that each paradigm has to offer. We then discuss KANT, a practical MT system which makes use of techniques from the interlingua, statistical, and transfer-based MT paradigms to produce accurate, high-quality translation.

1 Introduction: Burning Questions For MT

Should we pursue machine-aided human translation, human-guided machine translation or fully automated machine translation? Which is better: interlingua or transfer? Should MT be knowledge-based, grammar-based, case-based, or statistics-based? Should MT be applied to well-circumscribed domains, open-ended domains or should it be totally domain independent? Researchers often answer these questions more on the basis of their presupposition and training than rational argumentation or empirical confirmation. Perhaps the fault lies not with the typically inadequate answers provided, but rather with the questions themselves.

Consider, for instance, a well-trained carpenter being asked: Should carpentry use manual tools or human-guided power tools, or even a Japanese-style "lights-out" factory with full robotic automation? Should carpentry be a knowledge-based process, a strict follow-the-rules process, or a process based only on repeating past experience? Should carpentry be applied in narrow domains (e.g. kitchen cabinet making), or without regard to what is being built or repaired?[1]

The most glaring fallacy in the questions posed to and by MT researchers is the pressuposition of an exclusive disjunction among the alternatives. Does a carpenter use only a saw or only a hammer, when circumstances may dictate the most appropriate tool to be a screwdriver or a paint brush or a combination thereof? In the construction of even a single type of item some parts may be automatically mass-produced, others prepared using power tools, and some unique features handcrafted. Similarly in MT a combination of different techniques judiciously integrated may prove far more robust, economical, and capable of producing quality output than a single paradigm religiously followed for all texts, languages, styles, domains and end uses of the translation. The correct question then should be: which combination of methods can best do the specified MT job in the range of languages and domains necessary. "Best," of course, is defined by the end user, and usually implies a combination of translation accuracy, stylistic appropriateness, development and maintenance cost, integration into the operational environment, and economy of operation.

Another related fallacy is the assumption that all methods are known today, and we merely need to select one. New (albeit unproven) tools such as case-based translation (Sato, 1991; Sumita and Iida, 1991) emerge from time to time. Still another fallacy is the belief that once a translation method is proposed, it will never improve in fundamental ways, and it is sufficient to critique the early half-formed paradigm, rather than evaluate the current, more refined, and often far more effective technology. We would not evaluate carpentry tools and methods by looking only at their medieval counterparts; let us not do so to the more rapidly changing MT technology. An example of this problem is the mischaracterization of knowledge-based interlingual MT as impractical because of its voracious appetite for detailed, hand-coded knowledge, evident in the early KBMT systems – a condition much less true of the modern approach, as presented in the third section of this paper. Another example is seen in the burden of failures in the 1940's and 1950's weighing heavily upon statistically-based MT, whereas the present approach is at least in

[1] The reader who does not accept the close parallelism between the questions posed in this paragraph and those posed in the previous one is invited to draw a similar parallel between MT and computer system or user interface development. Should computer languages and interfaces be domain-oriented? How much automation of user functions should they assume? Etc.

some aspects far more promising (Brown, et al., 1990).

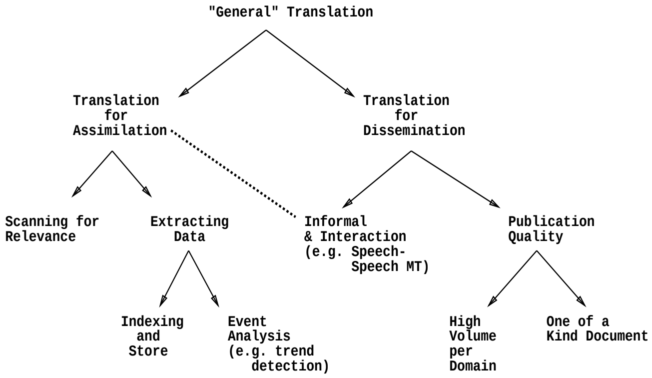

Figure 1: **Taxonomy of MT Tasks**

The last presupposition fallacy we discuss is that of referring to MT in the singular, rather than addressing each type of MT task in detail and not assuming that the best solution for one task is necessarily the best for all. Consider the taxonomy of potential MT tasks shown in Figure 1. The first major distinction is made between assimilation and dissemination tasks (Carbonell, et al., forthcoming). The former typically requires translation from many source languages into one's own language, and often the domains of translation are difficult to circumscribe. However, translation quality, especially for relevance scanning, need not be very high. Translation for dissemination, on the other hand, typically occurs between one's own language and multiple target languages, often in one or more predefined domains (e.g. operation and maintenance manuals for products), and requires high accuracy and fluency (i.e., mistakes in the operation or maintenance of jet aircraft engines are unacceptable). Exceptions and refinements to this generalization occur, however, leading to further taxonomic categorization, such as accounting for the event analysis task (under the assimilation category) separately, as it requires higher-quality translation than scanning and indexing. In brief, the domain and linguistic diversity demands placed on the MT technology are in large part defined by the type of MT required, and therefore paradigmatic questions are better raised in well-situated contexts where there is hope of arriving at more meaningful answers.

The remainder of this paper addresses each major MT paradigm and some possible combinations thereof in light of these better posed questions. Then, it focuses on KANT, an increasingly successful MT system based on a methodologically integrated approach for practical, high-accuracy MT in well-defined domains. Although primarily interlingual in nature, KANT draws on statistical methods, structural transfer ideas, and even some case-based MT to produce robust translations without postediting. The discussion combines rational reconstruction (what should work best for each task or stage in the translation) with empirical confirmation (what actually does work best in practice — through corpus analysis, experimentation and testing).

2 The Extreme Views

At the risk of caricature *in extremis*, let us consider some of the more radical views taken by proponents of particular MT paradigms. While many of these views have been voiced in workshops, MT conferences, and other discussions, some have also appeared in print. No explicit references are given here in order to protect the guilty (first author of this paper included), and to arrive quickly at the main point of the discussion.

2.1 Ye Olde Transfer

Transfer-based MT represents the incumbency, the entrenched paradigm, and with good reason; it led to the first successful applications of MT. Some proponents, especially in Britain, have argued that syntactic-transfer is correct in principle since they believe translation to be purely a surface linguistic transformation independent of meaning. Other proponents with a more practical turn of mind extend the paradigm to include semantic transfer, as discussed in the JTEC report (Carbonell, et al., forthcoming). The latter addresses some of the well-known problems with syntactic transfer such as partially redressing the lack of semantic disambiguation and ameliorating the N^2 problem for pair-wise translation among N languages by reducing the size of each transfer component. This latter direction goes part of the way to the integrated MT toolbox approach we advocate by incorporating some of the most easily usable tenets of interlingual MT.

Although the semantic transfer direction is a positive response to the shortcomings of the paradigm, other responses have been less forthcoming. Some (but far from all) transfer adherents, for instance, develop the standard old-guard response: "We were here first; therefore all others must be converted or squashed." Their favorite weapon is the attachment of derisory labels to all other paradigms. Hence, interlingual MT is "impossible;" KBMT is "impractical;" case-based translation is "unproven;" and statistical translation is "inaccurate." While there may be grains of truth behind these characterizations, the very process of attaching simplistic labels and closing the proverbial book on each new paradigm blocks scientific progress. MT practitioners should ask not what paradigms can be discarded but what ideas from other approaches can be borrowed, extended and integrated to improve the state of the technology.

2.2 Interlingua and the Kitchen Sink

Whereas interlingual MT and knowledge-based MT (KBMT) are not the same concept[2], in this section we

[2] Interlingual MT requires interpretation of the source text into a language-independent representation and subsequent generation into one or more target languages in order to eliminate the N^2 problem in multi-lingual systems. KBMT, on the other hand, implies the application of semantic and pragmatic knowledge – as well as syntactic and lexical knowledge – in order to produce more accurate and better translations. Fujitsu's ATLAS-II and NECs PIVOT, for example, are interlingua systems but are not truly knowledge-based. Likewise, one can

address the more extreme position where both concepts unite (Carbonell, Cullingford and Gershman, 1981; Nirenburg, Carbonell, Tomita and Goodman, 1991). The argument for interlingual KBMT is a simple one: Accurate translation requires comprehension. The main product of the comprehension is a semantic representation of the meaning of the source text, which can be exploited as the interlingua if designed in a sufficiently language-neutral manner. Ergo, in-depth understanding kills the two proverbial birds with one computational stone, the ambiguity resolution problem and eliminating the need for N^2 transfer grammars.

The obvious question, then, is why interlingual KBMT, being such a panacea to the ills of MT, not been universally embraced. The answer to this question is complex and informative, and very different from the old criticisms of interlingua (see Chapter 2 of (Nirenburg, Carbonell, Tomita and Goodman, 1991)).

First, there is an astronomical gap between the theoretical tenets of interlingual MT and their reduction to practice. In order for interlingua to be practical, large knowledge sources must be built: lexical, grammatical and semantic; full-scale system engineering is required of a somewhat greater complexity than for the traditional transfer paradigm. Perserverence pays off in MT. Principled and well-funded perseverence pays off even better. Only recently is interlingual KBMT being put to a true test as it is being developed for real applications. All previous argumentation was not situated in the context of measurable task performance.

Second, some interlingua partisans advocate a cancerous growth in the complexity of the interligua and of the knowledge sources. A semantic representation at the sentential level is not enough; one needs paragraph-level and text-level interconnected representations. Discourse analysis should be included with speaker-hearer attitudes thrown in for good measure. How about including indirect speech acts, metonomy, metaphor, implicit communicative goals, and the entire kitchen sink of lingusitic and dialog phenomena along with a massive proliferation in the number and complexity of supporting knowledge sources? And why stop here? Let us add common sense reasoning, the CYC knowledge-base, discourse planning for better generation, parallel processing and blackboards, the logic of belief-spaces and temporal relations, and so on, and on, and on.

Such is the attitude of some basic researchers, and it is good that scientists should explore the uncharted depths of full-scale reasoning and human knowledge. However, this is not the path to developing working machine translation systems in the foreseeable future. Instead, one must select what knowledge provides most of the action, following the well-known "80-20" rule. One should also see what simplifying assumptions can be made to produce very useful MT systems that may not (yet) address the full range of MT tasks shown in the earlier taxonomy.

easily conceive of a semantic-transfer-based KBMT approach not too different from the KANT system discussed in the following section, but without analyzing the source text into a true interlingua.

2.3 The Magic Wands of Statistics, Neural Nets and Case-Based Translation

Add one very large bilingual corpus, one sentence-to-sentence alignment process, a gallon of sophisticated Bayesian statistics, a touch of pixie dust, and throw them into the computational cauldron. Lo and behold, out comes a self-generated, robust, corpus-based, general-purpose machine translation system. Who needs a dictionary, grammars, semantics, or linguists? You can have it all, yes folks all of MT, using just the statistical method! This, of course, is a caricature of the statistical MT position. However, remove the pixie dust, and the caricature comes uncomfortably close to reality. Long before ALPAC, Warren Weaver (Weaver, 1949) believed that mathematical approaches could be as successful in MT as they proved to be in cryptology, but that line of inquiry led nowhere. The recent revival of statistical MT is better founded and has access to much more powerful computers, but has yet to prove its worth. Performance does not match SYSTRAN, and it takes minutes to translate a single short sentence on a mainframe. Therefore, statistical MT is far from the coveted silver bullet that will obviate the need for other more knowledge intensive techniques. Presently, adherents of statistical techniques are indeed recognizing benefits that can accrue through assimilation of other techniques (e.g. grammatical knowledge), and transfer and interlingua MT researchers are starting to find very useful roles for statistical analysis in their paradigms, as illustrated in the KANT discussion below.

Neural nets have been suggested as an answer to many of the World's problems, so why not use them for MT as well? Perhaps they may indeed prove useful in MT, as some preliminary work in JANUS suggests (Jain, et al., 1991), but they are essentially a kind of sophisticated statistical technique best suited to recognition and classification problems, such as low-level speech recognition.

Example-based translation (aka case-based MT) is becoming a serious alternative paradigm (Sato, 1991), but it is still in the early research phases. Adherents of this paradigm, fortunately, have not yet been fanatical in rejecting other paradigms, and as such may facilitate the integration of example-based and other translation methods. The KANT discussion below, in fact, shows one limited way in which example-based translation may play an important role at the phrasal level, albeit more in the role of example-based source-language analysis and example-based target language generation, rather than single-step translation.

In brief, none of the new paradigms promises a true revolution in MT. Instead, they promise to add novel tools to the MT toolchest, gradually producing a more robust and integrated technological infrastructure. For this reason, and only this reason, research into new translation paradigms is very useful.

2.4 Hackery Will Get You Nowhere

Though it may only take a mole to build a molehill, the process of building a mountain requires significantly greater resources. Likewise, the creation of a laboratory prototype capable of translating a few sentences cannot be equated with the construction of an operational MT system. As mentioned earlier, in MT perseverance counts. All of the successful systems, starting with the venerable SYSTRAN, have undergone a lengthy gestation period with significant resource commitment for system engineering and knowledge source (e.g., dictionaries, grammars) development. MT is not for the faint-hearted, nor is it for those of the hack-first-and-think-later persuasion. First of all, the process of developing the knowledge sources for full-scale translation outstrips the process of building the programs that use these knowledge sources. Second, developing multi-lingual and multi-domain translation require principled design and modularization (e.g., separation of language-specific data from language independent programs). Third, testing, maintaining and extending MT systems requires a well-developed methodology, version control and documentation. In short, MT systems require full-fledged software engineering, independent of MT paradigm(s). Quick-fix hackery is never the answer.

3 The KANT Perspective: A Rational Empiricist Approach

The KANT system represents a rational integration of previous approaches to machine translation. Although it is largely based on earlier knowledge-based, interlingua-style systems (e.g., KBMT-89), it also draws on important contributions from other translation methods and an overarching concern for practical applicability in real translation domains. In this section, we explore the relative roles of these methods in KANT. It is our goal to show that the success of the KANT system in practical domains is due to the internal use of different paradigms where they are most helpful in the translation process, rather than strict adherence to one particular paradigm.

KANT makes use of statistical methods (in corpus analysis, dictionary creation, and preference assignment), domain knowledge representation, and structural mapping rules (semantic interpretation and generation mapping) in order to achieve high-quality translation in the chosen domain. Thus KANT can be viewed as an amalgamation of crucial contributions from statistical translation, interlingual translation, and transfer-style translation.

3.1 The Role of Statistics

Any machine translation of practical utility must contain linguistic knowledge of thousands or tens of thousands of words and phrases in the source and target languages. High-quality translation depends in part on the depth and breadth of linguistic knowledge; a system which has more than a surface understanding of thousands of words is necessary to produce translations not requiring postediting. On the other hand, creating tens of thousands of lexical entries by hand is a time-consuming, labor-intensive process that can be prohibitively costly for some applications.

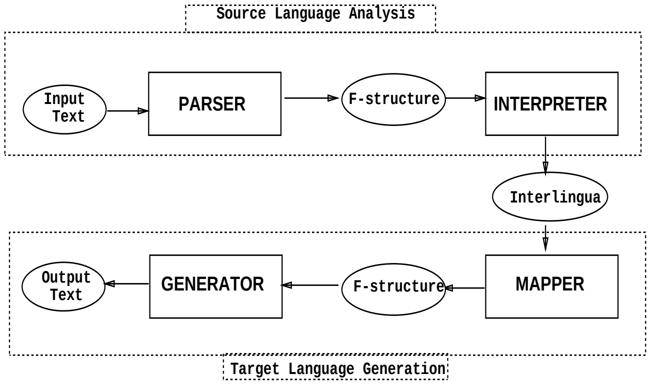

Figure 2: **Four Major KANT Modules**

Statistical analysis of source corpora can reduce dramatically the cost of linguistic knowledge acquistion in practical applications. Following computational analysis of source texts to identify potential words, phrases and their parts of speech, semi-automated programs can be utilized to create a complete pass at an initial lexicon. The subsequent refinement performed by the domain linguist is far more cost-effective, since the entire structure of the lexicon is in place and usually only minimal adjustment of the entry contents is necessary.

In addition, statistical analysis can be used to recover information about the particular senses of words used in the domain (for example, whether the term *gas* indicates gasoline or natural gas) and the structure of argument-taking words like verbs (i.e., whether or not a verb is transitive/intransitive, its subcategorization, etc.). For potentially ambiguous phrasal attachments (such as the ubiquitous prepositional phrase), statistics can be gathered which indicate a preferred attachment for phrases in certain contexts.

One of the strengths of case-based translation lies in the emphasis on finding phrases whose components have strong statistical collocation, and using these are the units to be translated. In KANT, statistical methods are used to deduce the set of phrasal terms (usually, technical terms expressed as complex noun phrases) that have integral meaning in the domain. These phrases can be mapped directly into a distinct interlingua concept, and do not require further internal analysis. Since the analysis of complex noun phrases is one of the harder problems in parsing, this can significantly reduce the chances for error.

All of these types of information are gathered by running statistical analyses on sample corpora (see Figure 3), or by the linguist inspecting domain examples using a KWIC indexing utility. The level of domain specificity that is achieved has a direct positive impact on the quality of translation, limiting the possible interpretation of words, phrases and sentences to just those that "make sense" in the domain, eliminating nonsense translations from the output text. Without using statistical tuning methods to limit these aspects of the lexicon, a system will admit many other ambiguous interpretations of the input which will lead to incorrect translation.

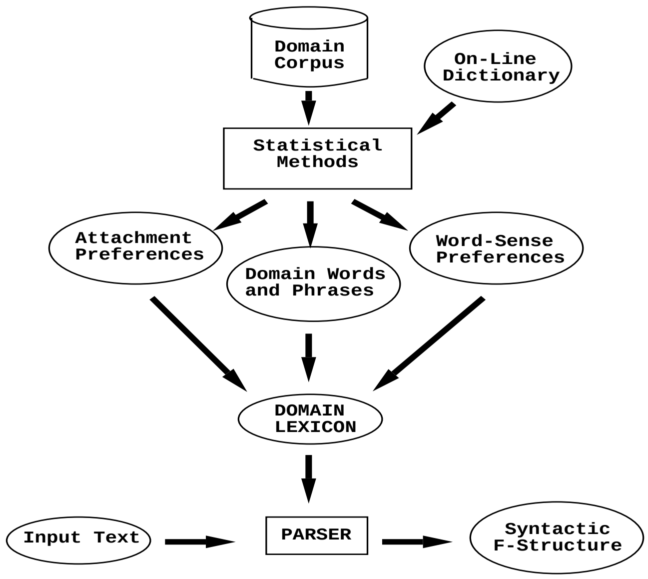

Figure 3: **Statistical Methods in Knowledge Acquisition**

3.2 The Minimalist Knowledge Base

Some research in computational linguistics has focussed on the exhaustive enumeration of an encyclopedic knowledge base for all world knowledge (Lenat and Guha, 1988). These efforts are interesting in the perspective of long-term research, since they raise many questions about the completeness of world knowledge and searching within a large knowledge base for concepts relevant to the current context, but they indicate that building a large knowledge base is quite costly and perhaps inappropriate for systems that must translate texts in a much more narrow domain. It is also the case that the proliferation of semantic features (corresponding to a highly-detailed level of representation for all aspects of meaning) which is necessary for stylistically excellent output (Nirenburg and Goodman, 1990) is not always practical or appropriate for the translation of matter-of-fact texts such as manuals, which require less elaborate prose.

In summary, KANT embodies the notion that practical translation requires that the knowledge base be limited to just what is necessary for stylistically adequate, accurate translation in a narrow domain. This in turn implies three basic objectives in knowledge base construction:

1. *Identification of Domain Concepts*.
   The domain model defines the semantic classes and instances required to represent all of the concepts present in the domain, but does not support any deeper reasoning in the domain. For example, if the texts to be translated make reference to automobiles but never to their subparts, then the knowledge base should not bother to represent a higher grain of detail (e.g., engines, tires, etc.). In a similar fashion, the modifying properties of concepts in the knowledge base should be limited to those actually mentioned in the domain; for example, if the color of an object is never referred to, then the knowledge base should not bother to represent color as a property.
2. *A Guide for the Construction of Interlingua Expressions*.
   Each concept definition includes not only a concept head (the "name" of the concept), but also a listing of its allowable semantic roles ("slots"), and allowable classes of concepts that the semantic roles can contain ("fillers"). By limiting the fillers of semantic roles to only certain classes of objects, events, or properties, the knowledge base defines what can and cannot be a valid interlingua expression following semantic interpretation.
3. *Disambiguation*.
   Role filler restrictions are crucial in the elimination of ambiguity that arises from multiple possible syntactic attachments (such as prepositional phrase attachment). In this way, the knowledge base defines what can and cannot be a valid interlingua expression following semantic interpretation, and thus helps to limit the process of source text analysis to only those interpretations possible within the domain. This is perhaps the single largest contribution of the knowledge base to the quality of the output.

The system's domain model is rich enough to allow all interpretations possible within the domain, but narrow enough to rule out irrelevant interpretations. The complexity of the domain model is only as deep as required to resolve ambiguity, which is the appropriate criterion for limiting the size of a domain model in a practical KBMT system.

3.3 Interpretation + Mapping = "Better" Transfer

Understandably, systems which are based on transfer between syntactic structures have much greater difficulty translating between languages that have very different syntactic structures (for example, between English and Japanese). Accurate, stylistically adequate output requires that the generation of the target text be based on the *meaning* of the source text and knowledge of the most adequate target lexemes and structures to express that meaning. This implies that an earlier process has somehow rendered the meaning of the original text in a form suitable for source-language-independent generation of a target text. These two steps are shown in Figure 2: the Interpreter and Mapper components of KANT.

This type of architecture alleviates the N^2 problem with pairwise transfer components. For each source language, the system must have a set of declarative rules of interpretation, which are compiled and used at run time to transform f-structures into interlingua expressions using the domain knowledge base as a guide. For each target language, the system must have a set of lexical selection and structural mapping rules, which select the appropriate lexical head and syntactic sub-structure (where appropriate) for each interlingua expression. As described in (Mitamura, Nyberg and Carbonell, 1991), this supports flexible lexical selection and text planning in generation, which are necessary for high-quality output.

For example, the English verb *overload* cannot be translated one-to-one into Japanese; an acceptable translation is *futan wo kake sugiru* ("excessively place a load"). During translation, *overload* is mapped into the interlingua concept `*e-overload`, which in turn maps into a noun phrase (*futan wo*), a main verb (*kake*) and an aspectual verb (*sugiru*). Thus the English sentence *Do not overload the outlet* translates to *konsento ni futan wo kake sugi nai de kudasai*, a sentence with quite different syntactic structure but identical meaning. In some cases, prepositional phrases in English must be translated to relative clauses in Japanese; for example, *I saw the man with a handkerchief* must be translated as follows:

```
Hankachi     wo  motta     otoko no hito   wo  mita
HANDKERCHIEF ACC HOLD+PAST MALE GEN PERSON ACC SEE+PAST
```

However, in other cases a *with* PP translates to a noun phrase, as in *I saw the man with a telescope* (instrumental reading):

```
Bouenkyou de   otoko no hito   wo  mita
TELESCOPE WITH MALE GEN PERSON ACC SEE+PAST
```

It is difficult to handle these distinctions without a complete interpretation of the input which takes into account semantic roles such as INSTRUMENT, POSSESSION, etc. Word- or phrase-based transfer breaks down when different words or phrases (like *with* or *with*-PP) must be treated differently on semantic grounds.

Even when translating between languages that are similar in their syntax (e.g., English and French), it is often necessary to select different words in the target language for the same source language word. For example, the English *reach* in *The temperature reaches the cloud point* must translate to the French *baisser* (*La température baisse au point de trouble*). However, *reach* must also translate to the French *monter* (*The temperature reaches the normal operating temperature* ⇒ *La température monte jusqu'à la température de marche normale*). The correct lexical selection cannot be made unless the system can represent the difference between a low and a high temperature. This is possible in a system with explicit semantics, and more difficult in a traditional transfer system.

In KANT, the single transfer process utilized in traditional systems is essentially replaced with two language-independent transfer modules which perform language-specific tasks based on declarative rules for each source and target language. Adding a new source or target language implies writing a new set of interpretation or mapping rules, without re-writing any of the system code, and connecting the new knowledge sources to the existing system for immediate translation to or from the new language.

The success of this architecture in producing high-quality output indicates that the process of transferring the syntax of one language to the syntax of the other should be mediated by a separate level of semantic representation. This in turn implies separate interpretation and mapping phases. As discussed above, these enhance the multilinguality and extensibility of an MT system, but it is certainly the case that rules must be written for dealing with the structural transformations from source to interlingua and from interlingua to target. In no way is interlingua seen as a "magic wand" that obviates the need for explicit structural rules for each language supported by the system. However, by splitting the transfer process into two discrete steps, knowledge acquisition becomes a language-by-language process rather than a pair-by-pair process, which is a much better way to transfer between the syntax of any two languages if the system is to support many languages. Moreover, having a non-ambiguous semantic interlingua as the "anchor point" alleviates problems such as lexical selection and structural disambiguation.

3.4 Results Thus Far

The KANT System is the result of approximately 10 man-years of effort, both in refining the basic software modules of the KBMT-89 system and in developing new software engineering techniques to make knowledge-based translation a practical possibility. In this section, we summarize the current status of the KANT system, its ongoing development in translation applications, and future directions.

3.4.1 Quality of Translation

KANT is capable of producing accurate translations of acceptable quality with no postediting (for examples, see (Mitamura, Nyberg and Carbonell, 1991). Of course, both accuracy and lack of postediting are achieved by extensive coding of mapping rules and lexical entries for each target language. Currently, KANT does not require any interaction with the author of the source text during translation. If the cost of an exhaustive domain knowledge base is prohibitive for a particular application, it is possible to incorporate a disambiguation step, where the system prompts the author to indicate which interpretation is the correct one (we have already investigated this technique in previous systems; for example, see (Brown, 1991)). Even in systems where interactive disambiguation of the source text is required, the cost is still far less than the cost of postediting in several target languages.

3.4.2 Practicality

Because KANT's initial applications are in the area of large-scale technical documentation, we have stressed practicality as a major motivating concern. Thus the space efficiency and speed of the code receive more emphasis than they would in a research-only system. In addition to the principles of minimal knowledge base construction, KANT must also employ a series of "knowledge compilers" that transform each knowledge source (grammar, lexicon, mapping rules, etc.) into a form that may no longer be human-readable but is far more efficient during computation. It is this step that is often lacking in typical prototype KBMT systems, where knowledge sources are utilized by the code in more or less the same format as they are created by the linguist. This knowledge compilation technology allows KANT to achieve translation speeds comparable to those of existing systems (e.g., 1-2 seconds per sentence), even when running on old workstations like the IBM RT PC (2.5 MIPS).

One current compromise that is made for the sake of practicality is the use of a source control language for text authoring. By creating a habitable source English that eliminates a few of the more problematic English constructions (from the point of view of computational translation), KANT is able to perform its task with stable, proven NLP technology. As parsing technology improves in its ability to handle difficult syntax, we hope to reduce source language control wherever possible.

3.4.3 Ongoing Development

The KANT system is currently being applied to the task of translating controlled English source documents (technical information manuals) to several target languages (including French, German, Spanish, Italian, and Japanese, with plans for others). The largest KANT application domain is service information manuals (operation/maintenance, assembly/disassembly, etc.) for heavy equipment products. We are also undertaking a preliminary investigation of multi-lingual translation for a large manufacturer of computer hardware and software, in the domain of operating system documentation.

The KANT prototype system (demonstrated at MT Summit III and described in (Mitamura, Nyberg and Carbonell, 1991)) shows great promise as the first practical refinement of the basic interlingual MT architecture. More comprehensive testing will be required to validate the KANT architecture and concomitant claims about robustness and quality in a larger-scale domain, but we will soon have the opportunity to test KANT in a much larger application domain as part of ongoing development.

3.4.4 Future Applications of KANT

As it exists today, KANT incorporates several years of development and refinement of basic NLP software tools developed at the CMT and utilized in a series of previous systems (CMT/SEMSYN, KBMT-89, Speech-Trans, etc.). Recent efforts have turned attention away from the creation of basic technology and more towards the support software that is necessary for large-scale practical application of MT systems. This support software includes statistical data gathering routines, interactive knowledge editor/browsers, compilers, etc. Perserverence and emphasis on support software has resulted in a basic architecture that not only performs well, but also supports quick, cost-effective development of new grammars, lexicons, etc. for new applications. In the future, it is our belief that the most commercially successful MT systems will support a high degree of modularity (separating knowledge from the system code itself) and application support software (for rapid knowledge acquisition and deployment in new domains).

4 Concluding Remarks

We must accept that translation is not for the faint-hearted, not for the quick-and-dirty approach, and not for the theoretical purist, but rather MT is for the long-haul pragmatist who recognizes the evolutionary nature of progress and does not suffer from the "not-invented-here" syndrome. Rational integration, empirical validation, an open mind, and a large dose of perseverance are the key ingredients.

In conclusion, we believe that continued evolution in MT research and development will prosper not by emphasizing paradigmatic differences between existing MT systems, but by integrating together successful techniques from each paradigm to create new systems that are better than the sum of their parts.

References

[1] Brown, P., J. Cocke, S. Della Pietra, V. Della Pietra, F. Jelinek, J. Lafferty, R. Mercer and P. Roossin (1990). "A Statistical Approach to Machine Translation," *Computational Linguistics* 16:79-85.

[2] Brown, R. (1991). "Automatic and Interactive Augmentation," in Goodman and Nirenburg, eds., *The KBMT Project: A Case Study in Knowledge-Based Machine Translation*, San Mateo, CA: Morgan Kaufmann.

[3] Carbonell, J., R. Cullingford and A. Gershman (1981). "Steps Towards Knowledge-Based Machine Translation, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 3:376-392.

[4] Carbonell, J., et al. (forthcoming). *The JTEC Report on Japanese Machine Translation*, National Science Foundation.

[5] Jain, A., A. McNair, A. Waibel, H. Saito, A. Hauptmann and J. Tebelskis (1991). "Connectionist and Symbolic Processing in Speech-to-Speech Translation: The JANUS System," *Proceedings of Machine Translation Summit III*, Washington, DC, July 2-4.

[6] Lenat, D. and R. V. Guha (1988). "The World According to CYC," MCC Technical Report (ACA-AI-300-88), Austin, TX.

[7] Mitamura, T., E. Nyberg and J. Carbonell (1991). "An Efficient Interlingua Translation System for Multi-lingual Document Production," *Proceedings of Machine Translation Summit III*, Washington, DC, July 2-4.

[8] Nirenburg, S. and K. Goodman (1990). "Treatment of Meaning in MT Systems," *Proceedings of the Third International Conference on Theoretical and Methodological Issues in Machien Translation of Natural Languages*, Austin, TX., pp. 171-188.

[9] Nirenburg, S., J. Carbonell, M. Tomita and K. Goodman (1991). *Machine Translation: A Knowledge-Based Approach*, San Mateo, CA: Morgan Kaufmann.

[10] Sato, S. (1991). "Example-Based Translation Approach to Machine Translation," *Proceedings of the International Workshop on Fundamental Research for the Future Generation of Natural Language Processing*, ATR Interpreting Telephony Research Laboratories.

[11] Sumita, E. and H. Iida (1991). "Experiments and Prospects of Example-Based Machine Translation," *Proceedings of the 29th Annual Meeting of the Association for Computational Linguistics*, University of California, Berkeley, June 18-21.

[12] Weaver, W. (1949). *Translation*, Reproduced in Locke and Booth, eds.: *Machine Translation of Languages*, Cambridge, MA: MIT Press, 1955.